

# Contents

<b>1</b>	<b>Partitioning Breaks Communities</b> . . . . .	<b>1</b>
	Fergal Reid, Aaron McDaid, Neil Hurley	
1.1	Introduction . . . . .	2
1.1.1	Cliques as Lower Bound Communities . . . . .	2
1.1.2	Partitioning Community Finding Algorithms . . . . .	3
1.1.3	Related Work . . . . .	3
1.2	Experiments . . . . .	5
1.2.1	Network Datasets Examined . . . . .	5
1.2.2	Partition by Modularity Maximisation . . . . .	6
1.2.3	Relation of Modularity Found to Proportion Split . .	8
1.2.4	Partition by Normalized Edge Cut . . . . .	9
1.3	Fundamental Partitionability of Networks . . . . .	10
1.3.1	Partitions that Directly Minimise Clique Splits . . . .	11
1.3.2	Detailed Analysis of Sample Networks . . . . .	12
1.3.3	‘Distinct’ Cliques . . . . .	13
1.3.4	Random and Synthetic Models of Community . . . . .	15
1.4	Overlapping Community Finding Algorithms . . . . .	17
1.4.1	Algorithms Examined . . . . .	18
1.4.2	Analysis in terms Of Split Cliques . . . . .	18
1.4.3	Community Overlap Graphs . . . . .	20
1.4.4	Analysis of Community Overlap Graphs of Overlapping CFAs . . . . .	22
1.5	Conclusion . . . . .	26
1.6	Further Work . . . . .	27
1.7	Acknowledgements . . . . .	27
	References . . . . .	27



# Chapter 1

## Partitioning Breaks Communities

Fergal Reid, Aaron McDaid, and Neil Hurley

### Abstract

Considering a clique as a conservative definition of community structure, we examine how *graph partitioning algorithms* interact with cliques. Many popular community-finding algorithms partition the entire graph into non-overlapping communities. We show that on a wide range of empirical networks, from different domains, significant numbers of cliques are split across the separate partitions produced by these algorithms. We then examine the largest connected component of the subgraph formed by retaining only edges in cliques, and apply partitioning strategies that explicitly minimise the number of cliques split. We further examine several modern overlapping community finding algorithms, in terms of the interaction between cliques and the communities they find, and in terms of the global overlap of the sets of communities they find. We conclude that, due to the connectedness of many networks, any community finding algorithm that produces partitions must fail to find at least some significant structures. Moreover, contrary to traditional intuition, in some empirical networks, strong ties and cliques frequently do cross community boundaries; much community structure is fundamentally overlapping and unpartitionable in nature.

**Key words:** Community Finding, Partitioning, Clustering, Network Analysis

---

Clique Research Cluster,  
Complex Adaptive Systems Laboratory,  
University College Dublin, Ireland  
e-mail: [fergal.reid@gmail.com](mailto:fergal.reid@gmail.com)

## 1.1 Introduction

Groups of interacting entities can be considered as a complex system. It is popular to examine such systems in terms of the networks their component entities form, to gain insight into properties of the system as a whole. For example, the speed with which a contagion can spread through a system is partly determined by the topology of its underlying network. The way sub-groups of entities interconnect is also important to investigate whether useful higher level abstractions – above the level of individual entities – exist in the systems we study. To find such structures, an extensive variety of algorithms have been developed, which attempt to find groups of nodes in the network that are structurally significant in some way; these groups are referred to in the literature as ‘communities’. See Fortunato [11] for an extensive review of these algorithms, which we will refer to as *Community Finding Algorithms*, or *CFAs*.

CFAs have been put to a range of applications, across several domains. As CFAs are applied ever more broadly, it is important that the structures they find, and the consequences of the design choices that define them are well understood. Particular CFAs should not be assumed to work across all complex networks, merely because they have evaluated well on some. In this research, we argue that certain algorithms, notably CFAs that produce *partitions* of the original network, return incomplete lists of the significant community structure, for at least some empirical networks. We perform an in-depth analysis of how several different CFAs interact with the cliques present in empirical networks, and discuss the consequences of this analysis for our intuition about community structure. We show that certain networks do not lend themselves well to partitioning, and caution against using partitioning algorithms as universal community finding tools.

### 1.1.1 Cliques as Lower Bound Communities

Each CFA finds structure that corresponds to a particular intuition of what a ‘community’ is; however there is little agreement on how exactly to define community. One common idea is that a community should have a high density of edges among its nodes, where *density* refers to the ratio of the number of actual edges between the nodes in the community to the maximum possible number of edges between these nodes. The bound of this definition is the graph theoretic structure known as a ‘clique’ – a fully connected subgraph, in which each node is connected to every other. Cliques, as discussed by Luce and Perry [22], have long been considered as community structure in human social networks. In the domain of social networks, this is particularly intuitive; if a user is friends with several others on Facebook, all of whom are also friends, then this is a significant structure of common friends. In

addition to this intuitive appeal, cliques are rare structures in the networks we study; due to the strict requirement for each node to connect to every other, clique structure is unlikely to arise by chance in a sparse network. Cliques are thus important structures. However, to define communities solely as cliques is very strict and conservative, for if even one connection in the group is missing – perhaps due to an incompletely observed network – then the found community will shrink. Many CFAs thus try to find communities comprised of groups of nodes which are highly connected, but less connected than perfect cliques. However, we posit that a clique is a good conservative *lower bound estimate* of community structure, in so far as an observed clique more than likely is wholly contained inside *some* real-world community. A maximally interconnected group of nodes, in a sparse network, always represents an interesting structure.

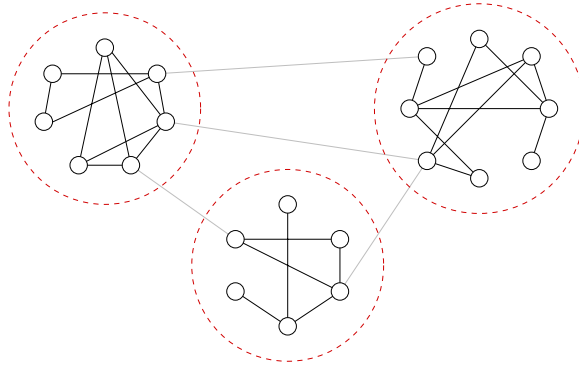
### 1.1.2 Partitioning Community Finding Algorithms

Many leading CFAs assign communities by partitioning the network, that is, grouping the nodes into *disjoint* subsets, assigning each node to exactly one subset. This partitioning approach to community finding has become popular, perhaps due to the appeal of treating a complex network as a graph, and the body of literature on graph partitioning problems. Early applications of graph partitioning, such as the applications of the Kernighan-Lin algorithm [16] discuss problems that explicitly require partitions, such as electronic component layout. However, in this work we are concerned about the completeness of the lists of community structure found by algorithms when used in other domains, such as social networks, and in complex networks generally. Regarding cliques as underestimates of community structure, we believe that regardless of what specific structures a given CFA finds, to be thorough, it should find, for each clique, *at least* one structure which is a superset of that clique. A CFA – considered as a tool that reveals structure in a complex network – that returns no community in which a group of fully connected nodes are assigned together, is neglecting to provide a complete list of the structures in the network. This is especially true if the clique is large in size.

### 1.1.3 Related Work

We show that in many complex networks, partitioning CFAs split cliques occurring within the network; and hence fail to find complete lists of the network structure. We examine why this occurs, investigating the intuition underlying many partitioning CFAs, and their relationship with cliques. We show, using cliques as a tool, that some traditional intuition describing com-

munities as well connected sets of nodes, separated by narrow bridges, is not always correct. Instead, many of the graphs we study exhibit a structure that can be better explained as the ‘pervasive overlap’ discussed in [1], [9] than as independent, weakly-connected modules. We analyse cliques, rather than any other community structure proposed in the ‘overlapping community finding’ literature, because we require a definition of structure that is a fundamental, conservative, and convincing underestimate of community; for every community, we want to find a conservative subset of that community. We use cliques, rather than structures such as the percolated  $k$ -cliques of Palla et al. [25], because with percolated  $k$ -cliques, we find no universal  $k$  consistent across networks with which to evaluate partitioning; this would make it difficult to be conservative in our analysis. Rather than choosing a new definition of community and discussing whether it is sufficiently conservative, we instead use the fundamental definition of the clique and examine its implications in detail. We analyse some of the same data as Leskovec et al. [21]. However, while that influential work sought to investigate the quality of the best community structure, at each scale, by evaluating it in terms of *conductance*, which penalises communities in proportion to their external edges, we instead investigate network structure from a different angle, by using the sociologically grounded idea of the clique to conservatively estimate community *cores*. We characterise to what level each and every clique is preserved after the network is partitioned, thus considering structures globally across the network.



**Fig. 1.1** Motivating image of network community structure from Newman [24]

An illustration of the intuition behind many CFAs can be seen in Figure 1.1, from the influential paper by Newman [24], which shows separate and well-defined modules, connected by only narrow bridges. This same intuition, conceptualising communities as connected by narrow bridges, can be traced back to the seminal work of Granovetter [13]: “*If the motivation to spread the rumor is dampened a bit on each wave of retelling, then the rumor moving through strong ties is much more likely to be limited to a few cliques than that going via weak ones; bridges will not be crossed.*”

Here, Granovetter is using ‘clique’ in the sociological sense, closer to the modern idea of community, and the idea is that bridges – narrow connecting links – need to be crossed to carry information between such cliques. This idea is further summed up in the modern review of Fortunato [11] as: *“If it were possible for a clique to move on a graph, in some way, it would probably get trapped inside its original community, as it could not cross the bottleneck formed by the inter-community edges.”* However, this work, in keeping with research [28, 30] on a limited number of other networks, finds evidence that structurally weak ties need not be crossed to traverse the network, contrary to the intuition just described. In fact, we show that while the traditional intuition may be appropriate in some cases, the structure of many empirical networks does indeed lead to cliques crossing the ‘bottleneck’ formed by inter-community edges.

## 1.2 Experiments

We conducted experiments to investigate the extent to which commonly used partitioning methods split the cliques in empirical network datasets. To keep the number of cliques we consider tractable, and in keeping with the original sociological definition of clique [22], we constrain our analysis to *maximal* cliques, which are cliques fully contained within no larger clique. For convenience, we refer to maximal cliques as simply ‘cliques’ in this work. In our analysis, we first generate the complete list of cliques present in each network using the fast Bron Kerbosch algorithm [3]. We then use the partitioning method under evaluation to assign each node to a community, and characterise how the cliques interact with the partitions found. We examine each maximal clique in turn, checking whether it is fully contained within a partition, or to what extent it has been split across partitions. We quantify and present this metric for each network, initially using two distinct partitioning methods; one popular and efficient modularity optimization method [2] and one normalised min-cut optimizing method [6].

### 1.2.1 Network Datasets Examined

To analyse data from a wide variety of networks, we gathered data from several different sources. We used several network datasets from the SNAP project<sup>1</sup> [21]. We examined networks formed by patterns of communication: The Enron and EU E-mail networks, and mobile telecoms data provided by

---

<sup>1</sup> <http://snap.stanford.edu/data/>

an industrial partner<sup>2</sup>, comprised of the voice call and SMS interactions on a mobile telecoms operator. We examined relation networks formed in online social networks, consisting of several Facebook university network datasets [29], samples taken from the full Twitter follower network [17], and the Slashdot online network. For both Twitter and the Mobile telecoms data, where we had access to very large networks, beyond reasonable computational means to analyse, we generated 3 random snowball samples of each network to produce tractable datasets. For the Facebook datasets, we chose to run our experiments on the smaller networks, due to the computational cost of calculating all maximal cliques. We also considered the SNAP academic publication networks, the Web networks of Stanford and NotreDame, product recommendation networks from Epinions and Amazon, and Wikipedia voting network. Finally, we considered a Protein-Protein interaction (PPI) network [5], as an example of a biological network.

### 1.2.2 Partition by Modularity Maximisation

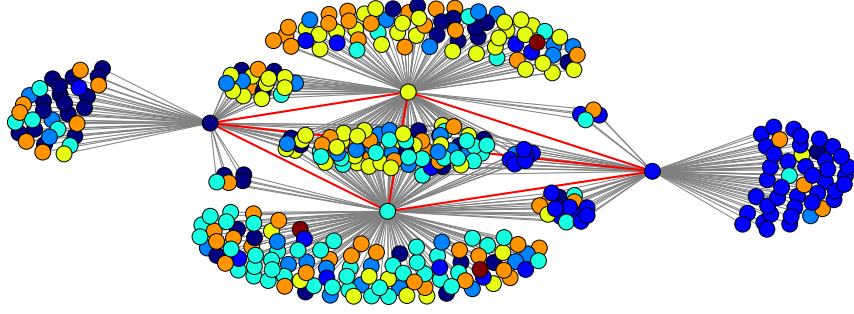
Many of the most popular CFAs are based on the modularity maximization approach of Newman [24]. The modularity function measures community quality as a count of internal edges, less the expected number in a random graph with the same node degrees. Modularity maximization algorithms, such as the fast ‘Louvain’ method of Blondel et al. [2] – which we evaluate here – designed to have a low computational cost on sparse graphs, and scale to large mobile call networks – optimise for the number of partitions as well as the associated partitioning. While traditional intuition holds that even triangles, or ‘strong ties’, should not cross community boundaries, we are interested in more significant cliques – so we initially restrict our analysis to cliques of size at least 4. We also use a conservative definition of when a clique is ‘split’ – we say a clique is “split at level  $\alpha$ ” if no partition contains more than  $(100 \times \alpha)\%$  of its nodes. We quantify the proportion of cliques that are split by the partitioning of each network in two ways. First, we examine the proportion of cliques of size at least 4 that are split at level  $\alpha = 0.9$ . Table 1.1 shows the significant proportions of cliques split at this level. We would have expected, based on the traditional intuition, that such structures would be contained in the center of the found communities – not spanning them, and not split by partitions that define found communities. Figure 1.2 provides an example of this effect, showing a single 4-clique that has been split across 4 separate partitions by the community finding algorithm.

As our metric is the proportion of maximal cliques that have been split, we might be concerned that many of the maximal cliques will be small, such as 4-cliques, and that if a 4-clique is split by partitions – while contrary

---

<sup>2</sup> Idiro Technologies





**Fig. 1.2** Visualisation of one of the split 4-cliques from the Caltech Facebook dataset. Clique edges are shown in red; modularity partitions, as found by the Louvain method, are shown by color; as can be seen, each node of this 4-clique has been assigned to a different community. This clique will thus not show up in the list of found communities. Note the many paths of length 2 between the clique’s nodes.

to the intuition of structurally strong ties being unable to cross community boundaries – this might not be of particular concern. For a more conservative experiment, we consider only large cliques of size at least 8, split at level  $\alpha = 0.8$ . These parameters are arbitrary and we do not seek to justify them other than to reiterate that we are considering conservative structure, which would traditionally be expected only in the ‘cores’ of found communities, not on their boundaries – structure that a comprehensive CFA should return. Even with this conservative definition, the partitions break significant numbers of such structures, on many networks – see Table 1.1. For example, this shows that the Louvain CFA, run on the Caltech Facebook network, will split over one quarter of cliques of size 8 or more.

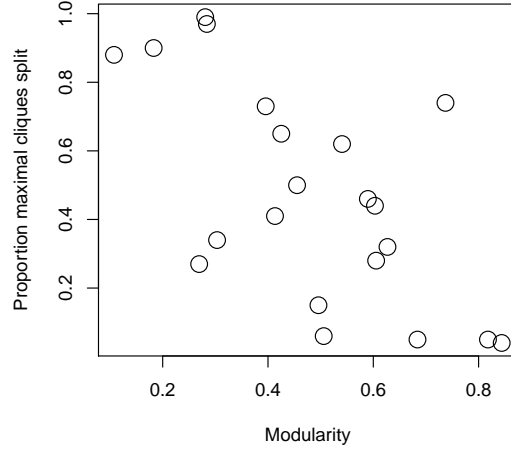
Our results show that the proportion of cliques split varies across the networks. There is also a large variation in the number of maximal cliques present. We might reason that this is due to some fundamental difference in the nature of the networks being considered, and question whether such analysis can be meaningfully applied across a range of networks. After all, the Amazon network is a network of frequently co-purchased products, and the web datasets are explicitly constructed lists of hyperlinks; still other networks involve human communication or collaboration. These networks are, however, frequently treated together as *complex networks*; we might *a priori* expect the same CFAs to perform well across them, and assume that a CFA proven in one domain will be automatically suitable and work well in other domains. However, this modularity method seems to do poorly on some types of network, at least where finding complete lists of community is desired. Similar results hold if we consider just the proportion of  $n$ -cliques split; as discussed in Section 1.3.2.

**Table 1.1** Proportion of maximal cliques split by the Louvain CFA, per network. We show the proportion of maximal cliques, of size 4 or greater, that have more than 10% of their nodes assigned to different partitions (i.e. are split at level  $\alpha=.9$ ). ‘Prop large cliques split’ is the proportion of maximal cliques, of size 8 or greater, that have more than 20% of their nodes assigned to different partitions (i.e. are split at level  $\alpha=.8$ ). ‘Cliques’ is the number of maximal cliques in the network, ‘Partitions’ is the number of partitions made by the Louvain method, and ‘Largest Clique’ is the size of the largest clique in the network.

Network	Nodes	Partitions (As found by Louvain Method)	Maximal Cliques	Largest Clique	Prop. Cliques split	Prop. Large cliques split
Email-Enron	36,692	1,363	205,712	20	0.61	0.47
Email-EuAll	265,009	15,743	93,267	16	0.82	0.67
Mobile1	10,001	182	1,550	10	0.97	0.00
Mobile2	10,001	124	3,538	10	0.90	0.00
Mobile3	10,001	86	951	9	0.88	0.00
Facebook-caltech	769	10	31,745	20	0.68	0.27
Facebook-princeton	6,596	21	1,286,678	34	0.44	0.22
Facebook-georgetown	9,414	26	1,440,853	33	0.41	0.22
Twitter1	2,001	8	23,570	12	0.99	0.66
Twitter2	2,001	4	554,489	27	0.15	0.01
Twitter3	2,001	7	130,399	22	0.06	0.00
Slashdot0811	77,360	771	441,941	26	0.13	0.01
Collab-AstroPhysics	18,771	331	27,997	57	0.60	0.32
Collab-CondMat	23,133	626	8,824	26	0.42	0.15
Collab-HighEnergy	9,875	483	2,636	32	0.23	0.00
Cite-HighEnergy	27,769	172	419,942	23	0.30	0.06
Amazon0302	262,111	173	117,054	7	0.01	0.00
Epinions	75,879	1,607	1,596,598	23	0.38	0.11
Web-NotreDame	325,729	693	130,965	155	0.04	0.00
Web-Stanford	281,903	1,013	774,555	61	0.04	0.01
Wiki-Vote	7,115	30	436,629	17	0.65	0.37
Protein-Collins	1,622	212	4,310	33	0.16	0.08

### 1.2.3 Relation of Modularity Found to Proportion Split

To investigate if the proportion of split cliques is in some way an artefact of low inherent modularity within the networks, we create a scatter-plot of the modularity achieved, against the proportion of maximal cliques split. From Figure 1.3 no obvious relationship appears. Several of the network partitions have high modularity and still display significant clique splitting; if there is a fundamental characteristic that renders particular networks unsuitable for modularity based partitioning, in terms of the proportion of cliques that will be split, then the modularity achieved does not capture it.



**Fig. 1.3** Scatter plot of modularity of the partition vs the proportion of maximal cliques  $>10\%$  split (i.e.  $\alpha = .9$ ), for each network.

#### 1.2.4 Partition by Normalized Edge Cut

Another method that has previously been used for the purpose of community finding, from a different family of algorithms, is the multilevel kernel  $k$ -means partitioning method implemented in *Grachus* [6], that minimises a normalized min-cut objective. Like the modularity maximisation method of Blondel et al. this implementation is designed to scale to large networks, performing well on sparse data by avoiding expensive eigenvector computation.

We examined this method on the same data as the modularity maximisation method. Unlike the modularity method, which discovers the number of partitions into which to break a graph, *Grachus* requires this to be specified. All other things equal, we would expect a smaller number of partitions would result in a smaller proportion of the maximal cliques being broken, and this effect is visible. However, even when asked to produce a relatively small number of partitions – relative to the network sizes – min-cut partitioning results in large proportions of the cliques greater than size 4 being split on many datasets, as shown in Table 1.2.

**Table 1.2** Proportion of cliques of size at least 4, split more than 10% (i.e.  $\alpha=0.9$ ), by *Grclus* [6], and *hMETIS* [15], per network. Values shown for 4, 16 and 64 Partitions, with *ufactor* 50, and 16 Partitions with *ufactor*500. Also shown, proportion of the large connected component preserved, for subgraphs of edges in at least 4-Cliques ‘4-Clique’, and edges in at least 5-Cliques ‘5-Clique’.

Network	Grclus			hMETIS			16	4-Clique 5-Clique	
	4	16	64	4	16	64	<i>uf</i> 500		
Email-Enron	0.38	0.74	0.92	0.10	0.54	0.67	0.38	0.55	0.39
Email-EuAll	0.53	0.86	0.98	0.20	0.58	0.76	0.42	0.04	0.02
Mobile1	0.75	0.88	0.99	0.47	0.81	0.93	0.80	0.17	0.07
Mobile2	0.66	0.93	0.97	0.47	0.77	0.92	0.64	0.20	0.09
Mobile3	0.83	0.93	0.96	0.48	0.82	0.95	0.77	0.06	0.01
Facebook-caltech	0.62	0.86	1.00	0.56	0.89	0.99	0.57	0.89	0.84
Facebook-princeton	0.33	0.69	0.89	0.32	0.58	0.89	0.36	0.92	0.89
Facebook-georgetown	0.30	0.58	0.80	0.32	0.50	0.74	0.40	0.93	0.90
Twitter1	0.88	0.99	1.00	0.82	0.97	1.00	0.83	0.78	0.57
Twitter2	0.22	0.99	1.00	0.05	0.88	1.00	0.56	0.70	0.57
Twitter3	0.74	0.98	1.00	0.04	0.65	0.99	0.05	0.74	0.33
Slashdot0811	0.28	0.49	0.94	0.08	0.13	0.37	0.09	0.10	0.04
Collab-AstroPhysics	0.43	0.53	0.77	0.27	0.49	0.65	0.34	0.83	0.71
Collab-CondMat	0.28	0.40	0.50	0.17	0.30	0.39	0.30	0.71	0.52
Collab-HighEnergy	0.16	0.28	0.43	0.10	0.19	0.29	0.19	0.42	0.13
Cite-HighEnergy	0.13	0.35	0.55	0.15	0.31	0.47	0.30	0.75	0.62
Amazon0302	0.01	0.02	0.04	0.00	0.00	0.00	0.00	0.11	0.00
Epinions	0.46	0.88	0.81	0.24	0.51	0.63	0.30	0.18	0.12
Web-NotreDame	0.01	0.03	0.11	0.00	0.05	0.18	0.04	0.07	0.03
Web-Stanford	0.03	0.04	0.39	0.00	0.09	0.46	0.02	0.49	0.40
Wiki-Vote	0.48	0.96	1.00	0.51	0.88	0.99	0.51	0.43	0.35
Protein-Collins	0.00	0.16	0.93	0.00	0.79	0.95	0.01	0.59	0.36

### 1.3 Fundamental Partitionability of Networks

Some datasets have a higher proportion of cliques split by partitions than others. This is largely uncorrelated with the mere number of cliques in the dataset, or the number of cliques per node, or per edge, or a number of other simple graph measures, such as clustering co-efficient. After investigating several popular CFAs, we now consider whether any partition exists which would not split cliques. Perhaps there were potential partitions that would confine cliques to the cores of the communities found, but these methods were not finding them? To answer this, we consider, for each network, the subgraph induced by nodes that share cliques; i.e. the network formed by discarding all edges from the network that are not part of cliques. The connected components in this subgraph are the sets of nodes that cannot be placed into separate partitions without splitting any cliques. We calculate the size of the largest connected component of each network, and present this as the proportion of nodes in the network, in Table 1.2.

We show results for the subgraph induced by nodes that share cliques of size 4 or greater, and of size 5 or greater, under the headings ‘4-Clique’ and ‘5-Clique’. On some networks, such as Facebook, Twitter, or collaboration networks, any partitioning scheme that is constrained to not split cliques of size five or greater has to leave the majority of nodes in a single partition.

This is an important structural property of these datasets, and an important result for certain diffusion models of complex contagion [4] which can only spread over structurally strong ties, as it shows these graphs are connected when using solely strong ties – it is possible to walk the graph communities without using weak ties. Further, on some of the larger datasets such as the Slashdot dataset, with 77,360 nodes, we find that over 30 per cent of those nodes (23,980 nodes) are in a connected component of the subgraph containing only edges that are in *triangles*; further evidence against the strict idea that strong ties do not cross community boundaries, and that communities are well separated.

### 1.3.1 Partitions that Directly Minimise Clique Splits

Having established the limits of partitions that break no *single* clique, we consider partitioning to directly optimise the number of cliques preserved, while producing balanced partitions. Partitioning a network while splitting as few cliques as possible is a hypergraph partitioning problem, where nodes in a clique together are connected by a hyperedge. This simple observation enables us to use a balanced mincut hypergraph partitioning algorithm, such as implemented by *hMETIS* [15] to partition the graph, while directly minimising clique splitting. *hMETIS* requires an important parameter to determine partition balance. Too high a value results in trivial partitions, with the vast majority of nodes in a single partition; too low might force *hMETIS* to make more aggressive hyperedge cuts than is reasonable. We initially set this *ufactor* at 50 (meaning the largest partition may have 50% larger weight than the average), to allow some unbalance. We examine cuts into 4, 16, and 64 partitions – generally fewer partitions than the modularity maximisation approach finds on these graphs. We also present results for 16 partitions with *ufactor* 500, allowing very large variation in partition size.

The results are shown in Table 1.2. Partitions directly minimising clique split indeed result in reduced proportions of the cliques split, compared to the balanced mincut of *Grclus*. As the number of partitions, and balance between partitions, constrain *hMETIS* more than the modularity maximisation method, the results are not directly comparable. However, as this method is directly minimising clique cut, it should approach a lower bound attainable by any partitioning CFA, for the given number of partitions – and, with generous balance parameters, indeed does better than modularity maximisation.

Even so, partitioning the network using this method, on a range of datasets – notably the collaboration networks, the Wiki voting data, the telecoms data and especially the Facebook social networks – still results in substantial proportions of cliques being split, demonstrating the fundamental global unpartitionability of some networks.

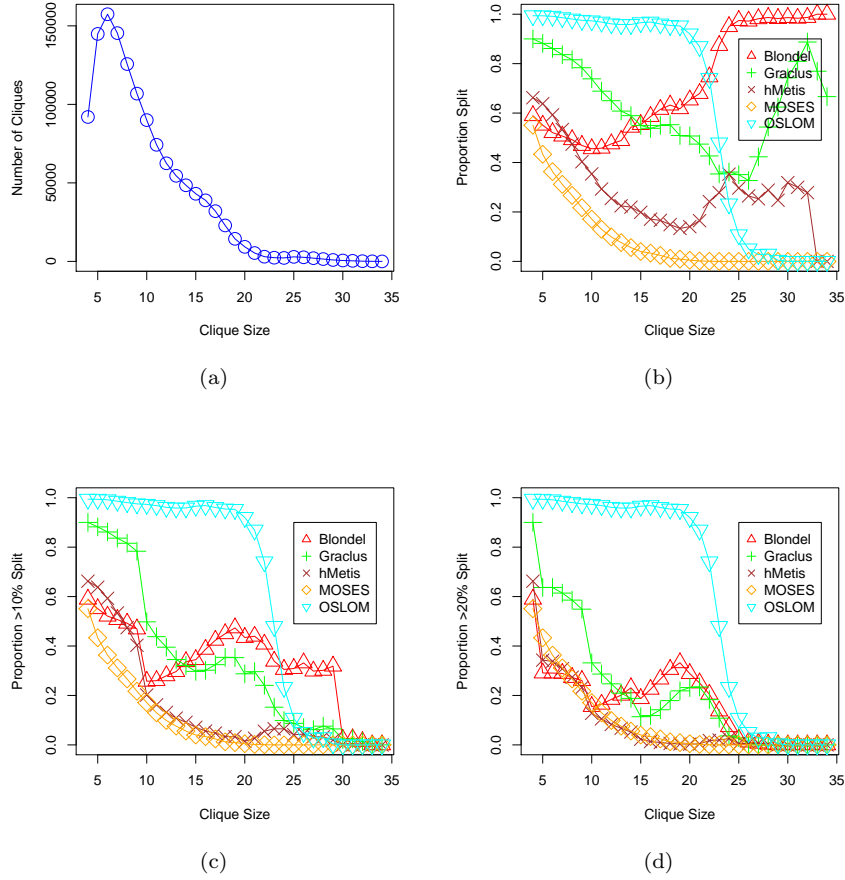
### 1.3.2 Detailed Analysis of Sample Networks

We now present some detailed statistics from three arbitrarily chosen sample networks: the Princeton Facebook network, which we will look at in detail as a case study, and one of each of the mobile and twitter sample networks. This Princeton Facebook network with over 6,500 nodes is large enough to allow us meaningfully investigate medium and large scale community structure. Facebook network data is also relatively dense, in that it captures many long term social relationships for each user; this is in contrast to more fleeting, or partial, network information we might obtain by extracting a network from a short term snapshot of a communications network.

In Figure 1.4 we show the number of cliques of each size in the network. We also show, for each clique size  $n$  observed in the network, the number of split cliques of that size. We plot this profile of cliques split, at each size, for each partitioning method investigated (as well as for non-partitioning Overlapping Community Finding Algorithms, which will be discussed in Section 1.4). We also show the proportion of cliques of size  $n$  split, for each value of  $n$ . We present results for three definitions of ‘split’ – where we consider cliques split if (b) any of their nodes have been partitioned from them, (c) greater than 10% of their nodes have been partitioned from them, and (d) greater than 20% of their nodes have been partitioned from them.

The Louvain method finds 21 partitions on this network; we use *Grclus* and *hMETIS* to produce the same number of partitions as the Louvain method. While the absolute number of cliques split tends to decrease as the metric becomes increasingly conservative, we note that in all cases, non-trivial numbers of cliques are split. As we would expect – as all partitioning algorithms try, in some sense, to avoid cutting edges – often the larger a clique is in size, the smaller the probability of the partitioning algorithms splitting it; however, cliques of all sizes are still split by these methods, on some networks; it is not the case that only the smallest cliques are split.

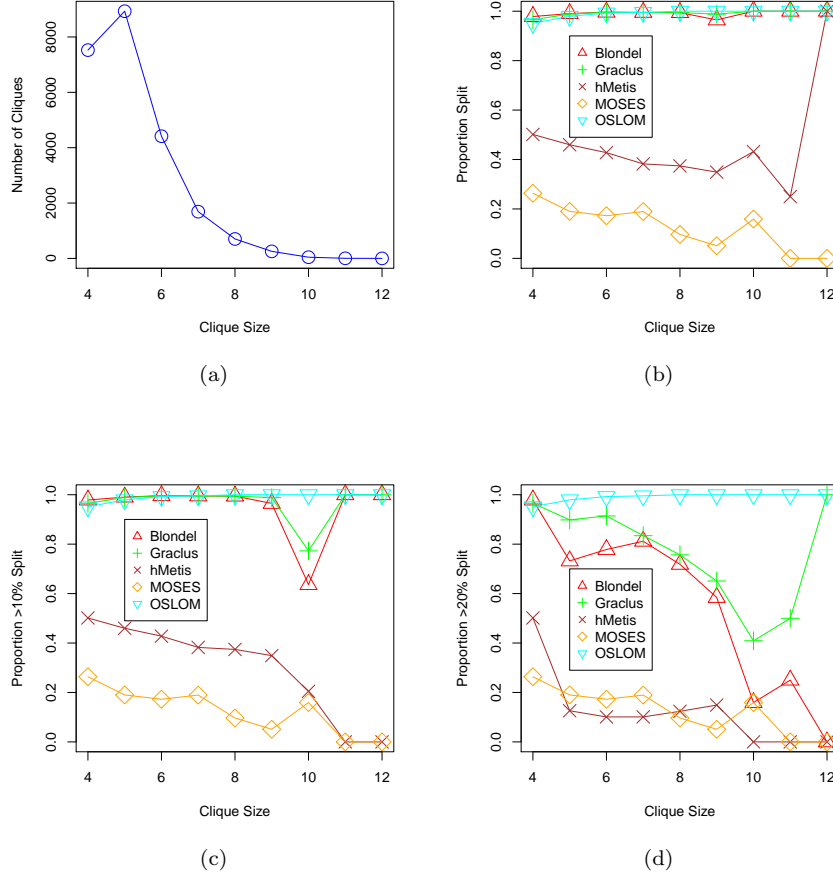
These figures emphasize the robustness of our findings – cliques of all sizes are split by partitioning – and illustrate an interesting way of characterising the effects of partitioning a network. Figures 1.5 and 1.6 show similar results for one of each of the Twitter and Mobile networks.



**Fig. 1.4** Proportion cliques split at each size, for Princeton Facebook Network (b) Considering a clique split if a single node is partitioned from it (c) Considering a clique split if  $>10\%$  of its nodes are partitioned from it (i.e.  $\alpha = .9$ ) (d) Considering a clique split if  $>20\%$  of its nodes are partitioned from it (i.e.  $\alpha = .8$ ).

### 1.3.3 ‘Distinct’ Cliques

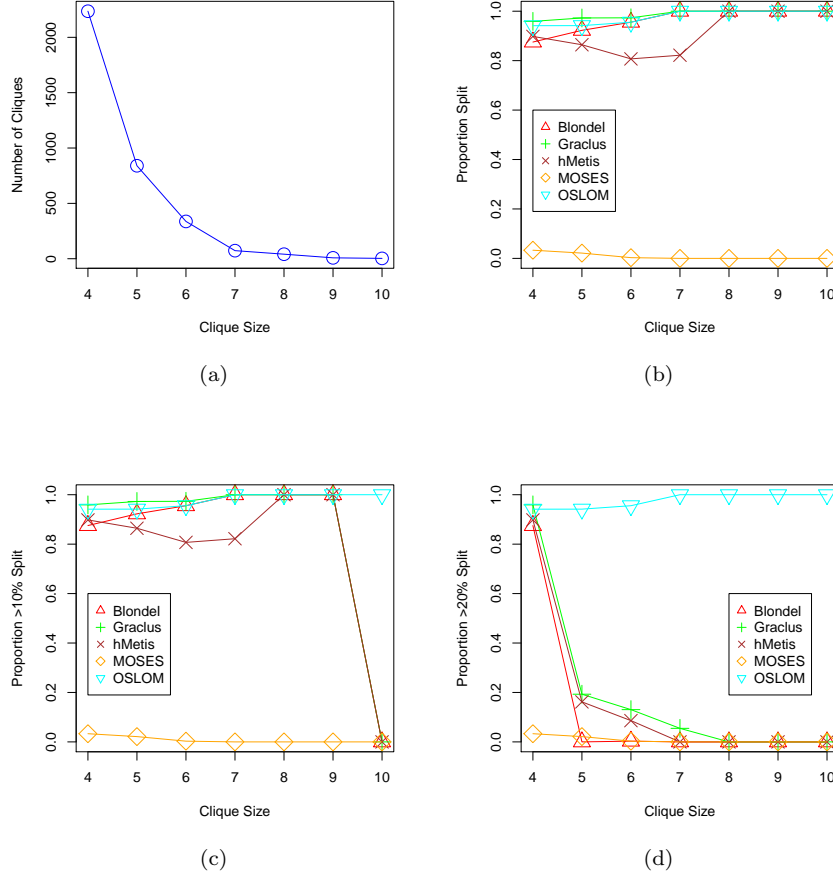
A large clique, with some small portion of random edges deleted, will turn into many very similar smaller cliques. In quantifying the ‘proportion’ of cliques split, we might be concerned that mis-assignment of a small set of nodes, if they are contained within a large number of very similar overlapping cliques, might skew the proportions. As an additional check on the robustness of these results, we present a set of results in Table 1.3 which correct for this effect



**Fig. 1.5** Proportion cliques split at each size, for one of the Twitter Networks (b) Considering a clique split if a single node is partitioned from it (c) Considering a clique split if  $>10\%$  of its nodes are partitioned from it (i.e.  $\alpha = .9$ ) (d) Considering a clique split if  $>20\%$  of its nodes are partitioned from it (i.e.  $\alpha = .8$ ).

by running our analysis not on the full set of maximal cliques, but instead on a set of maximal cliques, after a pre-processing phase which removes any clique that has a high Jaccard similarity ( $>0.8$ ) to any other larger clique. This analysis is computationally expensive to compute on the larger networks; however, on the networks we are able to perform it on, we find that our results still hold: substantial proportions of cliques are split, even if the only cliques we are looking at are cliques that are somewhat distinct from each other.





**Fig. 1.6** Proportion cliques split at each size, for one of the Mobile Networks (b) Considering a clique split if a single node is partitioned from it (c) Considering a clique split if  $>10\%$  of its nodes are partitioned from it (i.e.  $\alpha = .9$ ) (d) Considering a clique split if  $>20\%$  of its nodes are partitioned from it (i.e.  $\alpha = .8$ ).

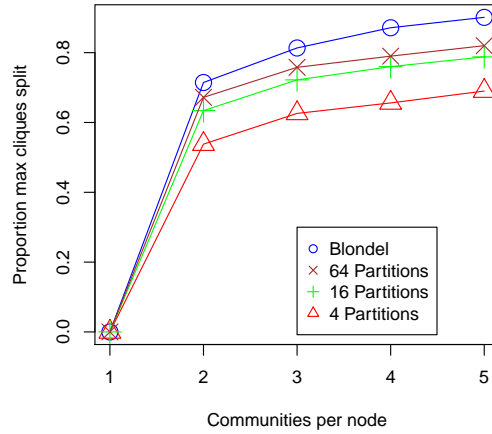
### 1.3.4 Random and Synthetic Models of Community

Broad categories of random community assignment model will produce networks where partitioning will fail to recover full communities. One source of synthetic benchmark community data is the ‘LFR’ benchmark [18], in which ‘communities’ – defined as sets of nodes with a high probability of edges between them – are embedded into a generated network. We ran our experiments on LFR graphs to test our method on synthetic data. We generated realisations of a 10,000 node network, altering the number of communities

**Table 1.3** Proportion of cliques that are ‘distinct’, beyond a given Jaccard similarity, that are over 10% split (i.e.  $\alpha=.9$ ) by *Grclus* [6], and *hMETIS* [15]. Values shown for *Grclus* and *hMETIS* for 16 partitions. *hMETIS* *ufactor* is 500.

Network	Louvain	Grclus	hMETIS
Email-Enron	0.62	0.76	0.39
Mobile1	0.97	0.88	0.80
Facebook-caltech	0.78	0.90	0.66
Twitter1	0.99	0.99	0.83
Collab-HighEnergy	0.23	0.28	0.19
Protein-Collins	0.34	0.33	0.02

each node was assigned to – from one to five, also increasing the corresponding number of edges, using the same parameters as with benchmarks detailed in previous work [20]. The results detailing the proportion of cliques split are shown in Figure 1.7.



**Fig. 1.7** Number of communities-per-node, as specified in benchmark parameters, vs proportion of maximal cliques  $>10\%$  split (i.e.  $\alpha=.9$ ), by the Louvain and hypergraph partitioning methods, on LFR benchmark data. Each data point is the mean of 5 LFR instances; deviation is negligible.

From these results, all methods partition the single-community-per-node networks without splitting cliques, but split significant numbers of cliques on networks with two or more communities-per-node. Even though the synthetic network model isn’t directly embedding cliques – just increasing edge density within communities – partitioning fails to find all structure, by our defined metrics, on synthetic networks where nodes are overlapping. Further, large components exist in the graph of edges in cliques in these generated networks. Not only are the individual nodes and communities overlapping as designed

by the model; it is a *global* property of the network as a whole that *no* non-trivial partition exists which does not split cliques.

## 1.4 Overlapping Community Finding Algorithms

A variety of algorithms exist which find overlapping communities within networks; Fortunato [11] mentions several of these, but this is an active area of research, with new algorithms frequently being developed [31]. Like with partitioning CFAs, many of these algorithms find subtly different structures, as authors work from slightly differing assumptions as to what constitutes a ‘good’ community. As such it is difficult to interpret what the output of a specific overlapping community finding algorithm tells us about the fundamental structure present in a network; and so we have avoided this in our analysis thus far. However, an advantage of overlapping CFAs is that they generally find structures that are much less common than maximal cliques are; typically a single overlapping community will contain many maximal cliques, many of which may differ only by a small number of nodes. As we have seen, a great number of cliques exist in the networks we examine; and while we can investigate aspects of network structure by using these cliques, the fact that so many of them exist brings some disadvantages; specifically, the ‘clique graph’ – the graph of cliques that overlap each other – is typically too large to work with.

In previous sections, we have considered the partitionability of networks, concentrating our analysis solely on cliques as the cores of community structure. The notion of cliques as community cores can be explicitly encoded in a community finding algorithm, both to produce communities that are disjoint, if disjoint cliques are enforced [32] or overlapping, if this criterion is relaxed [25], [27], [14]. Indeed, this is the approach of a family of overlapping CFAs, which use cliques as ‘seeds’ for communities, including the ‘Greedy Clique Expansion’ (GCE) algorithm [20], to which the authors of this work have contributed. The GCE method starts with maximal cliques as seeds and grows these seeds into communities using a local community quality measure. Thus, it will trivially produce communities in which there exists, for each maximal clique, at least one community that fully contains it. However, many other overlapping CFAs have kept with an ‘edge density’ notion of community quality, and find communities without any explicit modelling of cliques. Given the large numbers of cliques present in empirical networks, approaches that do not explicitly model cliques can have computational advantages over those that do. It is thus interesting to apply our clique based analysis to such algorithms and ask if density-driven community finding algorithms preserve clique cores, when communities are allowed overlap.

In this section, we will make use of overlapping CFAs for two separate purposes. First, we analyse two overlapping CFAs with the same procedure

as the partitioning CFAs, in order to ascertain the extent to which they split cliques in practice. Second, we examine the *community overlap graphs* created by these CFAs, and use the results of these graphs to examine the effects of partitioning on these networks, given the community structure as found by these particular overlapping CFAs.

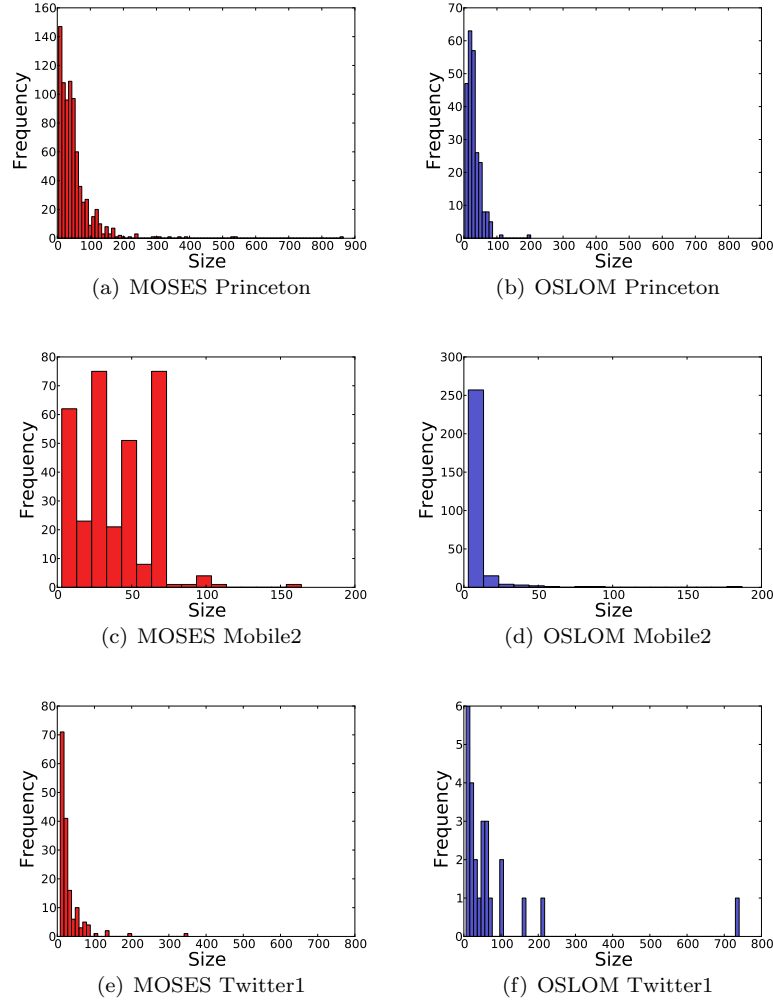
### 1.4.1 Algorithms Examined

We concentrate our analysis on two recent overlapping CFAs: MOSES (*Model-based Overlapping Seed Expansion*) which we proposed in [23] and OSLOM (*Order Statistics Local Optimization Method*) [19]. These statistically motivated algorithms do not explicitly use cliques in their operation, and are finding recent application in network analysis, for example the work of Grabowicz et al. [12]. While the complexity of these algorithms is dependent on the structure present in the input networks, like the previously discussed methods, both of these algorithms provide implementations which use heuristic techniques to enable them to scale to large networks, this makes them suitable for our analysis. We first examine the outputs of these algorithms. Figure 1.8 shows the community size distributions for each of the CFAs we analyze, on the case study networks we presented in detail earlier. We present distributions, rather than average community sizes, as the distributions of community sizes found vary widely by network, and tend to be heavily skewed. We do not include ‘singleton’ communities, containing only single nodes; OSLOM in particular reports many of these.

### 1.4.2 Analysis in terms Of Split Cliques

We analyzed the overlapping communities found by OSLOM and MOSES, in the same manner as the partitioning algorithms – for each clique, we check to see if there is any community in which it is fully contained; if there is not, we consider the clique to be split.

A thorough comparison of the exact structures found by these algorithms, each motivated by slightly differing models of community structure, is outside the scope of this work. To be thorough, we would have to either deal in subtle differences in the definition of ‘community’ – for which many definitions exist – or analyze the communities found by these methods against some ‘ground-truth’ data particular to a specific application domain. To restrict our analysis solely to network structure, we do not consider the issue of whether the communities found by MOSES and OSLOM are overall ‘good’ communities; instead we maintain our focus on split cliques. We present the results of this analysis in Table 1.4. We also show detailed results, for our case-study



**Fig. 1.8** Size distribution of overlapping communities found by MOSES and OSLOM. We do not show communities consisting of isolated nodes – OSLOM in particular finds a great many of these.

networks, on a per size-of-clique basis in Figures 1.4, 1.5 and 1.6. These results show that MOSES produces a set of communities such that most larger cliques, in most networks, are contained in at least one community found by MOSES. This is an interesting result, considering that MOSES does not explicitly find communities in terms of cliques. The benchmarking of OSLOM yields a different result, however: for large numbers of cliques, OSLOM does not produce at least one community containing the clique.

It is difficult to explain this. Unlike MOSES, OSLOM outputs a hierarchy of levels of community, and we only consider the lowest level of that hierarchy; perhaps, in practice, the lowest levels are very ‘fine grained’ for OSLOM, below the level of an individual clique. Alternatively, in any community finding algorithm, there must always be a tradeoff between the sensitivity required to find all communities, and specificity to avoid finding ‘false positive’ communities. We can use cliques as underestimates of community structure, to measure sensitivity – in that every clique should be contained in a community – but not to measure specificity, for as discussed earlier, it may be too strict to require that every community contain a clique. Perhaps OSLOM is simply more specific in its output than MOSES; the two algorithms find structures of different quantity and size, as Figure 1.8 shows. A detailed investigation of these issues would have to be undertaken in the context of a specific application domain, with ground truth data. But what these results do show is that at least some overlapping CFAs, which contain no explicit representation of cliques, find communities which split much fewer cliques than the partitioning algorithms do. The results also show that while the use of an overlapping CFA is necessary to avoid splitting cliques, as discussed in Section 1.3 and concretely illustrated by the results of MOSES, it is not a sufficient condition in practice as the OSLOM results show. Thus if a specific application domain requires high sensitivity, and a full list of community structure to be found, then not only must an overlapping CFA be used, but the CFA should also be evaluated for the specific application.

### 1.4.3 *Community Overlap Graphs*

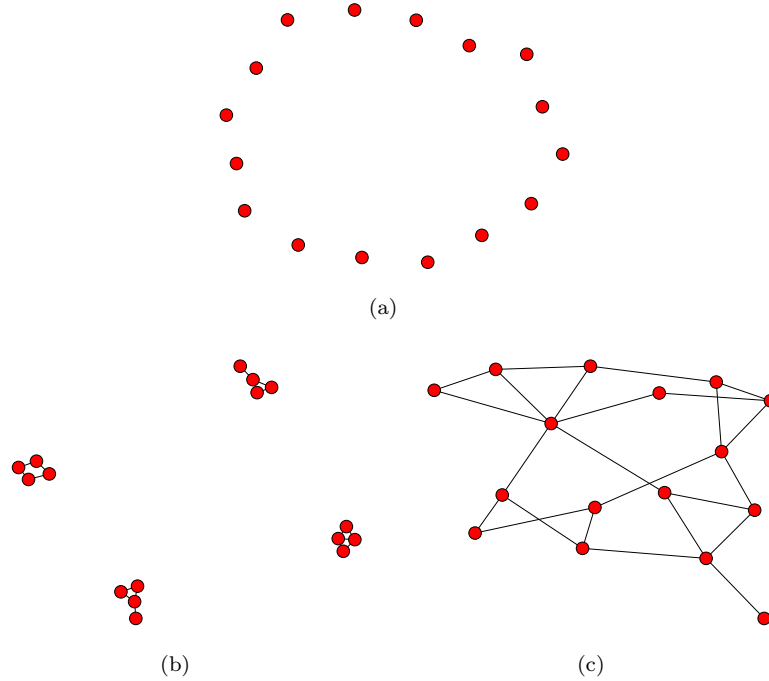
We have considered, in Section 1.3, the fundamental partitionability of networks, by examining the connected components which exist when we only consider subsets of the networks connected by cliques. We have also considered whether a hypergraph partitioning method, attempting to split as few cliques as possible, can partition the network, and have seen that – assuming cliques at the core of communities – communities overlap each other pervasively. We can develop our intuition about these results further, by considering the possibilities for the structure present in the ‘*Community Overlap Graph*’ (COG); the graph formed by representing each community as a node, and connecting a pair of nodes (communities) by an edge, when the two communities overlap by more than some threshold number of nodes i.e. when they share more than some threshold number of nodes in common. When considering individual cliques as our communities, this idea is identical to that of the ‘clique graph’ discussed by Everett and Borgatti [10] and discussed further by Evans [8].

However, the large numbers of maximal cliques present in the networks we study make explicitly working with the clique graph difficult. The commu-

**Table 1.4** Proportion of cliques that are not completely contained in at least one community – i.e. are ‘split’ – by the OSLOM and MOSES overlapping CFAs. Some networks present in previous benchmarks are not shown, due to the algorithms taking too long to run.

Network	OSLOM >10% split	OSLOM Size >8 >20% split	MOSES >10% split	MOSES Size >8 >20% split
Email-Enron	0.96	0.98	0.16	0.01
Email-EuAll	0.93	0.00	0.06	0.00
Mobile1	0.99	0.00	0.00	0.00
Mobile2	0.94	0.00	0.03	0.00
Mobile3	0.99	0.00	0.03	0.00
Facebook-caltech	0.76	0.37	0.41	0.04
Facebook-princeton	0.93	0.76	0.21	0.01
Facebook-georgetown	0.95	0.82	0.22	0.01
Twitter1	0.97	0.94	0.20	0.00
Twitter2	0.98	0.91	0.02	0.00
Twitter3	0.98	0.91	0.02	0.00
Collab-AstroPhysics	0.86	0.79	0.35	0.04
Collab-CondensedMatter	0.40	0.22	0.08	0.02
Collab-HighEnergy	0.33	0.00	0.08	0.00
Cite-HighEnergy	0.81	0.52	0.15	0.00
Amazon0302	0.09	0.00	0.01	0.00
Epinions	0.96	0.83	0.01	0.00
Wiki-Vote	0.79	0.67	0.04	0.00
Protein-Collins	0.13	0.06	0.06	0.01

nities found by overlapping CFAs however, are typically smaller in number (as shown in Tables 1.5 and 1.6). Different possibilities for what structure we might see in the community overlap graph are shown in Figure 1.9. We can see that Figure 1.9(a) corresponds to a world view of non overlapping communities, in which the partitioning of networks into communities makes obvious sense. Figure 1.9(b) contains overlapping communities, but, perhaps surprisingly, it still makes some sense to partition the network, with partitions dividing clusters of overlapping communities together. We have shown, from our analysis of paths through cliques, and attempting to partition the network using hypergraph partitioning on the found cliques, that a world-view similar to Figure 1.9(c) is most appropriate. We will now discuss these ideas in more detail, with reference to actual community overlap graphs, generated from the communities found by the two overlapping community finding algorithms, on empirical data.

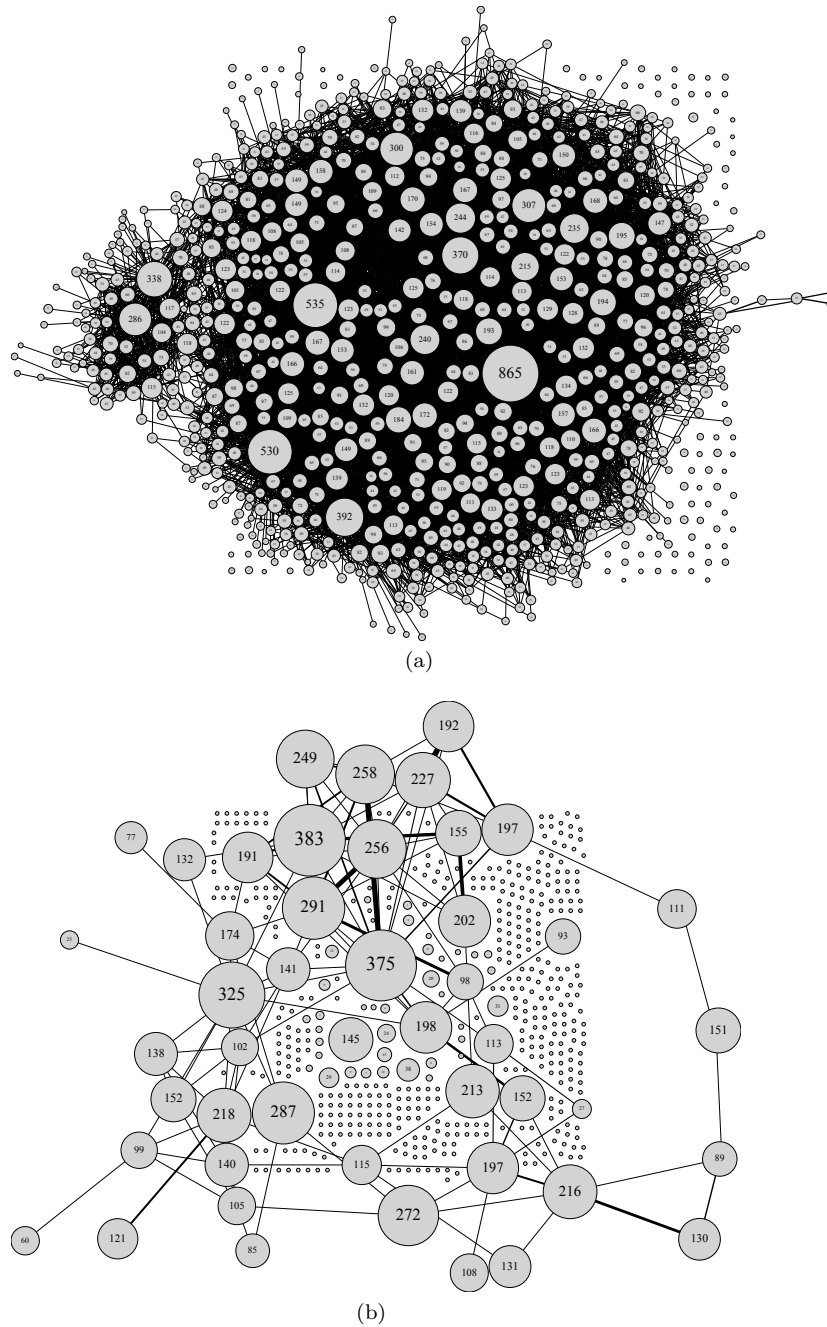


**Fig. 1.9** Illustration of some possibilities for the Community Overlap Graph, for a network with 16 communities. Each node represents a community; edges connect pairs of overlapping communities. (a) Non-overlapping communities. (b) Overlapping communities, but clustered, with no path through overlap. (c) Overlapping communities with unpartitionable overlap.

#### 1.4.4 Analysis of Community Overlap Graphs of Overlapping CFAs

In Figure 1.10 we show visualisations of the Community Overlap Graph of the communities found by the MOSES and OSLOM algorithms, in the Facebook Princeton network. In order to show only the more significant overlaps, we draw an edge between two communities overlapping by at least 4 nodes. These visualisations show that in this particular network, most of the larger communities found by these two algorithms – and hence most of the nodes in the network – are part of a connected component of overlapping communities. As such, the empirical visualization corresponds most closely to Figure 1.9(c), and hence partitioning to find communities is not suitable in networks like this. Visualizing the community overlap graphs of these networks shows clearly the extent to which communities overlap, and the structure that would be broken by partitioning these networks in order to find communities.



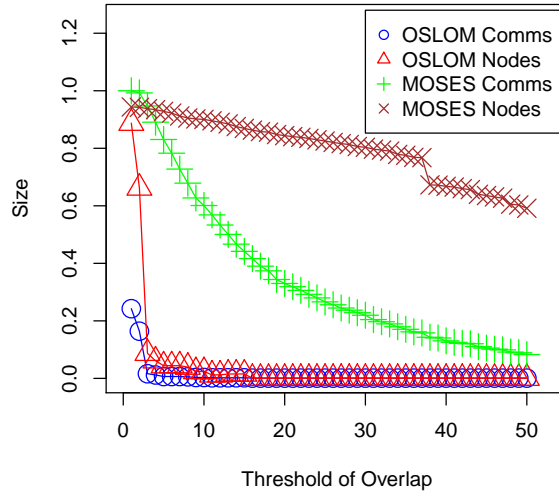


**Fig. 1.10** Visualisation of COG of Facebook Princeton network. An edge is drawn whenever two communities overlap by at least 4 nodes. Edge width is proportional to overlap, and node area is proportional community size. Communities are labeled with the number of nodes they contain. Nodes may be present in multiple communities; two communities with a high degree of overlap contain fewer unique nodes than the sum of their labels. Shown here is the COG extracted from running (a) MOSES and (b) OSLOM on the Facebook Princeton network. Networks visualised with Graphviz [7] force directed layout.

In addition to visualizing these networks, we can attempt to quantify the degree to which a set of overlapping communities is partitionable, similar to how we examined the fundamental partitionability of networks in Section 1.3. We attempt to quantify this by examining how many of the communities in the community overlap graph are in the largest connected component of that graph, and what proportion of the nodes in the source network they contain. If a large proportion of communities and nodes are in a connected component, then this again would indicate quantitatively that the community structure is closer to Figure 1.9(c) than (a) or (b).

Our results for the MOSES method are presented in Table 1.5, and for OSLOM in Table 1.6. As we can see, in line with our earlier results using cliques, the degree of overlap varies across networks with the social networks – particularly the Facebook networks – showing the least partitionable results.

In line with the results obtained by quantifying the proportion of cliques split, MOSES finds structure that is more highly overlapping than OSLOM. These results show an interesting method of quantifying the degree of overlap of community structure in a given network, and for a given overlapping CFA.



**Fig. 1.11** As the threshold of overlap is changed, the size of the largest connected component in the community overlap graph changes. We investigate how the size of this component varies both in terms of the number of communities in it, and the total number of nodes connected by communities that overlap by at least that threshold. We display the size of the largest component, both in terms of the proportion of communities that are in it, and in terms of the proportion of nodes in the underlying graph which are in it, for both OSLOM and MOSES Overlapping CFAs

**Table 1.5** Results for the size of the largest connected component (LCC) of the community overlap graph (COG), examining community structure found by MOSES.

Network	Nodes	Number Comm- unities	Comms in LCC of COG	Nodes in LCC of COG	Proportion Nodes in LCC COG
Email-Enron	36,692	2,471	587	14,573	0.4
Email-EuAll	265,009	473	257	24,919	0.09
Mobile1	10,001	437	38	1,159	0.12
Mobile2	10,001	323	219	8,609	0.86
Mobile3	10,001	171	120	8,478	0.85
Facebook-caltech	769	81	71	666	0.87
Facebook-princeton	6,596	797	710	6,162	0.93
Facebook-georgetown	9,414	893	800	8,740	0.93
Twitter1	2,001	161	132	1,686	0.84
Twitter2	2,001	129	99	1,680	0.84
Twitter3	2,001	188	101	1,080	0.54
Collab-AstroPhysics	18,771	2,816	677	9,953	0.53
Collab-CondMat	23,133	3,458	175	3,760	0.16
Collab-HighEnergy	9,875	1,663	15	274	0.03
Cite-HighEnergy	27,769	1,625	998	21,445	0.77
Amazon0302	262,111	23,665	47	1,767	0.01
Epinions	75,879	795	249	19,889	0.26
Wiki-Vote	7,115	65	63	3,805	0.53
Protein-Collins	1,622	150	16	221	0.14

In Tables 1.5 and 1.6 we used an overlap of 4 nodes as a threshold for ‘significant’ overlap between two communities. It is interesting to examine how the threshold used to analyze the community overlap graph effects the connectivity of that graph. Figure 1.11 shows how the threshold effects the size of the largest connected component in the community overlap graph, both in terms of the number of communities in it, and the number of nodes of the underlying network that are in it, for both of the overlapping CFAs we examined. It can be seen from this figure that OSLOM has a sharp falloff in the size of the LCC as the threshold of overlap is increased. MOSES exhibits a much more gradual falloff in the size of the LCC – even if we require that communities have 10 nodes in common for them to be overlapping, the graph is still largely unpartitionable, without breaking several community overlaps. This difference between MOSES and OSLOM is perhaps not surprising, given MOSES’s tendency to find larger communities than OSLOM, and is consistent with the results in terms of the proportion of cliques split. This is further evidence that the level of overlap in the structures found by varying overlapping CFAs can vary widely.

**Table 1.6** Results for the size of the largest connected component (LCC) of the community overlap graph (COG), examining community structure found by OSLOM.

Network	Nodes	Number Comm- unities	Comms in LCC of COG	Nodes in LCC of COG	Proportion Nodes in LCC COG
Email-Enron	36,692	10,620	46	2,722	0.07
Email-EuAll	265,009	131,143	-	-	-
Mobile1	10,001	8,435	1	3	0
Mobile2	10,001	8,119	4	235	0.02
Mobile3	10,001	9,195	2	157	0.02
Facebook-caltech	769	137	2	100	0.13
Facebook-princeton	6,596	920	11	404	0.06
Facebook-georgetown	9,414	1,189	5	178	0.02
Twitter1	2,001	467	21	1,529	0.76
Twitter2	2,001	113	7	1,463	0.73
Twitter3	2,001	289	9	1,040	0.52
Collab-AstroPhysics	18,771	4,106	9	202	0.01
Collab-CondMat	23,133	5,911	6	159	0.01
Collab-HighEnergy	9,875	3,808	7	145	0.01
Cite-HighEnergy	27,769	4,393	12	358	0.01
Amazon0302	262,111	37,374	19	424	0
Epinions	75,879	46,260	57	3,739	0.05
Wiki-Vote	7,115	2,744	21	4,085	0.57
Protein-Collins	1,622	529	2	50	0.03

## 1.5 Conclusion

We have investigated a wide range of empirical networks, characterising them according to the proportion of cliques in them that are split by various partitioning methods. Our results show that the early intuition on how communities are embedded in graphs does not hold across all networks and domains. On many complex networks cliques do not exist *solely* in community cores connected only by narrow bridges and weak ties – instead they frequently overlap across the community boundaries produced by partitioning algorithms.

If we accept cliques as conservative lower bounds for community structure, then, on many networks, partitioning CFAs are fundamentally limited in the completeness of the communities they can find, as shown by our results on the graph of edges in cliques, and from using hypergraph partitioning algorithms to partition cliques. This shows that communities are not easily separable from each other simply by removing structurally weak ties; instead, communities overlap across each other, with pairs of community frequently connected by strong ties, and other communities.

Our analysis of overlapping community finding algorithms has shown that some overlapping CFAs produce sets of communities in which each individual clique is fully contained. However, as we have shown, not all overlapping CFAs satisfy this property. We have presented community overlap graphs as

another tool – in addition to cliques – with which to explore the output of overlapping CFAs. We have shown that on some networks, the communities as output by overlapping CFAs reveal a community structure that cannot be partitioned.

Thus, caution is warranted when using partitioning community finding algorithms where there is a sensitivity requirement that all significant community structure be found. In agreement with recent research on pervasive overlap, conceptualising networks as overlapping meshes of strong ties, with denser community regions, and using a CFA designed to find communities that overlap, will be more appropriate in many application domains.

## 1.6 Further Work

Work on formal models of community generation that might explain whether a network is suitable for partitioning, and attempt to characterise the generative processes behind this global overlap would be interesting. That cliques frequently span communities also has implications for the type of diffusion processes that can occur on networks; data on the non-partitionable nature of communities may lead to an enhanced understanding of diffusion on complex networks. We hope that studying the nature of community overlap can lead to a better fundamental understanding of structure in empirical networks, and help development of future community finding algorithms.

## 1.7 Acknowledgements

This work is supported by Science Foundation Ireland under grant 08/SRC/I1407: Clique: Graph and Network Analysis Cluster. An earlier version of this work appeared in *ASONAM '11* [26].

## References

1. Y. Ahn, J. Bagrow, and S. Lehmann. Link communities reveal multiscale complexity in networks. *Nature*, 466(7307):761–764, 2010.
2. V. Blondel, J. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008:P10008, 2008.
3. C. Bron and J. Kerbosch. Finding all cliques of an undirected graph. *Communications of the ACM*, 16(9):575–577, 1973.
4. D. Centola and M. Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–34, 2007.

5. S. Collins, P. Kemmeren, X. Zhao, J. Greenblatt, F. Spencer, F. Holstege, J. Weissman, and N. Krogan. Toward a comprehensive atlas of the physical interactome of *Saccharomyces cerevisiae*. *Molecular & Cellular Proteomics*, 6(3):439, 2007.
6. I. Dhillon, Y. Guan, and B. Kulis. Weighted graph cuts without eigenvectors: a multilevel approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1944–1957, 2007.
7. J. Ellson, E. Gansner, E. Koutsofios, S. North, and G. Woodhull. Graphviz and dynagraph static and dynamic graph drawing tools. *Graph Drawing Software*, 2004.
8. T. Evans. Clique graphs and overlapping communities. *Journal of Statistical Mechanics: Theory and Experiment*, 2010:P12037, 2010.
9. T. Evans and R. Lambiotte. Line graphs, link partitions, and overlapping communities. *Physical Review E*, 80(1):016105, 2009.
10. M. Everett and S. Borgatti. Analyzing clique overlap. *Connections*, 21(1):49–61, 1998.
11. S. Fortunato. Community detection in graphs. *Physics Reports*, 486:75–174, 2010.
12. P. Grabowicz, J. Ramasco, E. Moro, J. Pujol, and V. Eguiluz. Social features of online networks: The strength of intermediary ties in online social media. *PloS one*, 7(1):e29358, 2012.
13. M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360, 1973.
14. F. Havemann, M. Heinz, A. Struck, and J. Gläser. Identification of overlapping communities by locally calculating community-changing resolution levels. *Arxiv preprint arXiv:1008.1004*, 2010.
15. G. Karypis and V. Kumar. Multilevel k-way hypergraph partitioning. In *Proceedings of the 36th annual ACM/IEEE Design Automation Conference*. ACM, 1999.
16. B. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, 49(2):291–307, 1970.
17. H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.
18. A. Lancichinetti and S. Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E*, 80(1):016118, 2009.
19. A. Lancichinetti, F. Radicchi, J. Ramasco, and S. Fortunato. Finding statistically significant communities in networks. *PloS one*, 6(4):e18961, 2011.
20. C. Lee, F. Reid, A. McDaid, and N. Hurley. Detecting highly overlapping community structure by greedy clique expansion. *KDD SNA 2010*, 2010.
21. J. Leskovec, K. Lang, A. Dasgupta, and M. Mahoney. Statistical properties of community structure in large social and information networks. In *Proceeding of the 17th international conference on World Wide Web*, pages 695–704. ACM, 2008.
22. R. Luce and A. Perry. A method of matrix analysis of group structure. *Psychometrika*, 14(2):95–116, 1949.
23. A. McDaid and N. Hurley. Detecting highly overlapping communities with model-based overlapping seed expansion. In *Advances in Social Networks Analysis and Mining (ASONAM), 2010 International Conference on*, pages 112–119. IEEE, 2010.
24. M. Newman. Detecting community structure in networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 38(2):321–330, 2004.
25. G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.
26. F. Reid, A. McDaid, and N. Hurley. Partitioning Breaks Communities. *2011 Advances in Social Network Analysis and Mining*, 2011.
27. H. Shen, X. Cheng, K. Cai, and M. Hu. Detect overlapping and hierarchical community structure in networks. *Physica A: Statistical Mechanics and its Applications*, 388(8):1706–1712, 2009.

28. X. Shi, L. Adamic, and M. Strauss. Networks of strong ties. *Physica A: Statistical Mechanics and its Applications*, 378(1):33–47, 2007.
29. A. Traud, E. Kelsic, P. Mucha, and M. Porter. Comparing community structure to characteristics in online collegiate social networks. *SIAM review*, 53(3):526–543, 2011.
30. M. van der Leij and S. Goyal. *Strong ties in a small world*. Tinbergen Institute, 2006.
31. J. Xie, S. Kelley, and B. Szymanski. Overlapping community detection in networks: the state of the art and comparative study. *Arxiv preprint arXiv:1110.5813*, 2011.
32. B. Yan and S. Gregory. Detecting communities in networks by merging cliques. In *Intelligent Computing and Intelligent Systems, 2009. ICIS 2009. IEEE International Conference on*, volume 1, pages 832–836. Ieee, 2009.